

Optimisation of Simulations of Stochastic Processes by Removal of Opposing Reactions

Fabian Spill,^{1,2} Philip K. Maini,³ and Helen Byrne³

¹*Department of Biomedical Engineering, Boston University, 44 Cummington Street, Boston MA 02215, USA*

²*Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*

³*Wolfson Centre for Mathematical Biology, Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK*

(Dated: 9 February 2016)

Models invoking the chemical master equation are used in many areas of science, and, hence, their simulation is of interest to many researchers. The complexity of the problems at hand often requires considerable computational power, so a large number of algorithms have been developed to speed up simulations. However, a drawback of many of these algorithms is that their implementation is more complicated than, for instance, the Gillespie algorithm, which is widely used to simulate the chemical master equation, and can be implemented with a few lines of code. Here, we present an algorithm which does not modify the way in which the master equation is solved, but instead modifies the transition rates, and can thus be implemented with a few lines of code. It works for all models in which reversible reactions occur by replacing such reversible reactions with effective net reactions. Examples of such systems include reaction-diffusion systems, in which diffusion is modelled by a random walk. The random movement of particles between neighbouring sites is then replaced with a net random flux. Furthermore, as we modify the transition rates of the model, rather than its implementation on a computer, our method can be combined with existing algorithms that were designed to speed up simulations of the stochastic master equation. By focusing on some specific models, we show how our algorithm can significantly speed up model simulations while maintaining essential features of the original model.

I. INTRODUCTION

Stochastic models are used in many areas of physics, chemistry and biology to describe random fluctuations associated with uncertainties in the external environment and the intrinsic discreteness of the species of interest, which could be individual atoms, molecules, cells or animals ^{1,2}. The chemical master equation is commonly used to determine the time evolution of the probability density function which describes the current state of the system of interest. Such a state could, for instance, be defined by the current number of molecules of each kind in a model of chemical reactions ³⁻⁵, the number of animals of each species in an ecological model ⁶, or the number of cells in an animal tissue ⁷⁻⁹. The problem which remains is then to solve for the time-evolution of the probability density function.

As analytic solutions to the master equations are rarely obtainable, one typically aims to simplify the master equation, (e.g. by the van Kampen expansion ¹), or one numerically computes solutions of the master equation. A commonly used algorithm for such simulations is the Gillespie algorithm ^{10,11}. It is exact in the sense that every reaction is taken into account, and no stochastic information is lost in the simulation process. However, this makes it computationally expensive, as for each event several potentially expensive computational steps, such as the simulation of random numbers, the recalculation of propensities or the identification of the next occurring reaction need to be performed. Additionally, in many cases, the large number of elementary reactions involved makes the use of the conventional Gillespie algorithm practically unfeasible.

For this purpose, a number of algorithms have been developed in order to speed up the simulation of the master equation. Some of these are hybrid methods that switch from a stochastic to the corresponding mean field model ¹²⁻¹⁶ in regions of space or time where the latter are reasonable approximations to the stochastic equations. Others, such as the tau-leap method ¹⁷⁻²², simulate several reactions within one time-step. Related work focuses on more efficient simulation of stochastic models not based on the chemical master equation ²³⁻³².

A common feature of the aforementioned methods is that their implementation is typically more complex than for the Gillespie algorithm, which can be implemented with a few lines of code. In this paper, we introduce an algorithm which does not affect the solution method of the master equation, but rather changes the transition rates of the master equation itself. As

a result, existing algorithms (and their optimisations) can be used, and only the definition of the transition rates needs to be changed. The idea is to group together reversible reactions. For example, a particle which is created and subsequently annihilated before it reacts with any other particle has no observable effect. Likewise, in a stochastic diffusion problem on a lattice, if two particles of the same type swap sites, this has no observable effect on the state of the system, provided no reactions happen while the particles switch locations. Our algorithm involves combining transition rates for reactions which have opposing effects to produce effective rates that describe the net effect of the two reactions. The model, with the modified transition rates, is designed such that the mean behaviour is the same as that of the original stochastic model, but typically stochastic effects such as variance are reduced. Hence, we suggest that the modified transition rates are used when the number of particles involved in the reaction exceeds a threshold value chosen such that stochastic effects corresponding to the reaction of the transition rate which we modify can be ignored. This is similar to methods such as ¹⁶ which replace parts of a stochastic model with the mean-field description of the same model. In the mean-field limit of a stochastic model, the dynamics of two opposing reactions is also only taken into account in a net, effective way. However, unlike the mean-field limit, our method is based entirely on the chemical master equation. Hence, it is easier to implement than most other hybrid methods, as the modifications to our algorithm are based on the transition rates and, hence, can be easily combined with transition rates of other reactions which are not optimised. Furthermore, in contrast to other methods^{16,28}, which approximate a stochastic process in some regions by its mean-field description, our method does not encounter technical challenges at the interface of the stochastic and mean field domains.

The remainder of the paper is organised as follows. We discuss a simple birth-death process in section II. This example serves to introduce our algorithm, and can have applications in complex reaction networks where at least some reactions are reversible. In section III we focus on the stochastic diffusion of particles randomly migrating on a lattice. This is a good test of our algorithm since when two particles move in opposite directions, the system returns to its original state. We show that our model of stochastic diffusion accurately preserves stochastic properties such as first-passage times. Our study of stochastic diffusion is also important for understanding how our algorithm will perform for the more general case of reaction-diffusion systems. For instance, when diffusion occurs on a faster

timescale relative to reactions, many diffusive events may occur which do not affect the reactions, but slow down simulations of the system. In section IV A we focus on the stochastic Fisher-Kolmogorov system to how that our algorithm can increase simulation speed significantly while maintaining essential stochastic features, such as the modification of the wave speed due to stochastic effects. In section IV B, we study a one-dimensional model of Min oscillations. This multi-species example illustrates how our algorithm performs when some molecular species are present at low copy numbers, leading to stochastic effects, while others are present in high numbers, making simulations slow. Again, our method can markedly improve simulation speed, while preserving the statistical distributions of species with low copy numbers.

II. CREATION AND ANNIHILATION PROCESS

We consider a spatially averaged model in which particles of type A can either divide, or annihilate when they hit another particle of the same type:



Here, λ , and μ are rate constants associated with division and annihilation, i.e. for small times t , $\lambda t + \mathcal{O}(t^2)$ is the probability that a randomly chosen particle divides, and $2\mu t + \mathcal{O}(t^2)$ is the probability that one of two randomly chosen particles annihilates the other. Such a model is the simplest to which our algorithm can be applied, as it consists of only two reactions which have equal and opposite effects. Our algorithm replaces the creation and annihilation reactions with a net reaction which, depending on whether creation or annihilation is more likely, will itself be a creation or annihilation reaction with a rate given by the difference of the rates of the original model.

We model this process with the chemical master equation, denoting by $P(N, t)$ the probability density function for N particles to be present in the system at time t :

$$\frac{dP(N, t)}{dt} = ((E^- - 1)\mathcal{T}_{N+1|N} + (E^{+1} - 1)\mathcal{T}_{N-1|N}) P(N, t), \quad N \in \mathbb{N} \tag{2}$$

Here, E^\pm are shift operators that shift the number of particles N by ± 1 , and $\mathcal{T}_{N+1|N}$ ($\mathcal{T}_{N-1|N}$)

denote the transition rates for the creation (annihilation) of a particle, so that

$$\begin{aligned}\mathcal{T}_{N+1|N} &= \lambda N, \\ \mathcal{T}_{N-1|N} &= \mu N(N-1).\end{aligned}\tag{3}$$

We note that $\mathcal{T}_{0|1} = \mathcal{T}_{-1|0} = 0$, so the model ensures that neither extinction nor negative particle numbers can occur, if N is initially positive. From equations (2), (3) we can see that if $\lambda N > \mu N(N-1)$, then a new particle is more likely to appear than to disappear. Conversely, if $\lambda N < \mu N(N-1)$, particles will be more likely to disappear. We thus define a second stochastic process with the same master equation (2), but transition rates given by

$$\begin{aligned}\mathcal{T}_{N+1|N} &= \max(\lambda N - \mu N(N-1), 0), \\ \mathcal{T}_{N-1|N} &= \max(\mu N(N-1) - \lambda N, 0).\end{aligned}\tag{4}$$

It is straightforward to show that, if $\overline{N^2} \approx \overline{N}^2$, which can be motivated, for instance, by the van-Kampen expansion ¹ in the limit of large particle numbers, then the model with transition rates (4) leads to the mean-field equation

$$\frac{d\overline{N}}{dt} \approx \lambda \overline{N} - \mu \overline{N}(\overline{N} - 1), \quad \overline{N} = \sum_N N P(N, t),\tag{5}$$

which is identical to that for the model with transition rates (3). However, the behaviours of (3) and (4) for low particle numbers and low ratios of $\frac{\lambda}{\mu}$ are quite different (see also Figure 1 and the discussion below). We note that for a given state, transition rates (4) are always always bounded above by those in (3). Therefore, a global change in state will typically involve fewer stochastic events in (4) than in (3) and be faster to simulate with the Gillespie algorithm. We henceforth refer to (4) as the difference model, as it is constructed by taking the differences of reversible reactions in the original, exact model (3). Having established that for large particle numbers (4) and (3) behave similarly, and that for small particle numbers they do not, we now introduce a model which conditionally switches between transition rates (3) and (4):

$$\begin{aligned}\mathcal{T}_{N+1|N} &= \begin{cases} \max(\lambda N - \mu N(N-1), 0) & \text{if } N \geq \Theta, \\ \lambda N & \text{otherwise,} \end{cases}, \\ \mathcal{T}_{N-1|N} &= \begin{cases} \max(\mu N(N-1) - \lambda N, 0) & \text{if } N \geq \Theta \\ \mu N(N-1) & \text{otherwise.} \end{cases}\end{aligned}\tag{6}$$

Here, we have introduced a threshold Θ such that when $N < \Theta$ we use the exact model (3), and when $N \geq \Theta$ we approximate (3) by the difference model (4). We will refer to (6) as the conditional difference model. If $\Theta \rightarrow \infty$ we recover (3), whereas (4) is obtained for $\Theta = 0$. Note that while (6) leads to the same mean-field equation as (3), its stochastic properties are, generally, different. We can view (6) as an approximation to (3), and the challenge is to choose Θ so that simulation results are in close agreement with (3). Figure 1 shows a




FIG. 1: Birth-death process with parameters $\lambda = 5$, $\mu = 1$ and $N = 20$ at $t = 0$, simulated with the original stochastic model (3) (obtained by setting the threshold $\Theta = 100$), the difference model (4) ($\Theta = 0$) and the conditional difference model (6) with $\Theta = 10$, and the mean-field ODE (5).

typical simulation of the birth-death process with parameters $\lambda = 5$, $\mu = 1$ and $N = 20$ at $t = 0$. We compare results generated from the exact model (3), the difference model (4), the conditional difference model (6) with a threshold value $\Theta = 10$ and the corresponding ODE model, equation (5). We see that the behaviour of the conditional difference model is qualitatively similar to that of the original model. By contrast, the number of particles in the difference model decays randomly until it reaches $N = \frac{\lambda}{\mu} + 1 = 6$, which is an absorbing state in the difference model. Hence, the difference model behaves qualitatively differently from the original, exact model and the conditional difference model when $N \approx 6$. It behaves more like the mean-field ODE solution because of the absorbing state, which is not present in the original model (3), and when $N > 6$, particles in the difference model can only decay,

i.e. cannot be created. In supplementary Figure S1, we show the distribution of numbers for two thresholds, $\Theta = 0, 10$, at two times, shortly after the start of the simulations, $t = 0.01$, and at the long time, $t = 10$, when the solution is close to its steady state. In each case we compare the results to those of the exact model, which is obtained by fixing $\Theta = 100$. We note that the distributions at $t = 0.01$ are similar for all cases, but when the threshold is less than the initial condition, the distribution is bounded above by the initial condition. At $t = 10$, we notice that for $\Theta = 0$, the distribution approximates the delta distribution, whereas when $\Theta = 10$, it closely matches that for the exact model. We explain those results by noting that in the exact model, particle numbers exceed $N = 10$ only with low probability.

A more detailed comparison between the exact model and its approximation by the difference and conditional difference models will be performed in the sections that follow, where we consider the examples of diffusion and reaction-diffusion problems.

III. STOCHASTIC DIFFUSION MODELS

We now consider stochastic versions of the one-dimensional diffusion equation,

$$\frac{\partial n}{\partial t} = D \frac{\partial^2 n}{\partial x^2}, \quad (7)$$

defined for $x \in [0, L]$ and $t \geq 0$, and which is supplemented with boundary conditions which we do not specify here for generality. In equation (7), $D > 0$ is the constant diffusion coefficient, and the concentration $n = n(x, t)$ denotes the number of particles per unit length. We keep all quantities in dimensional form to better understand the physics. We discretise the spatial domain into k_{max} equally spaced compartments of size h . We denote by $N_k(t)$ the number of particles in compartment k at time t , dropping the time dependence when no confusion is possible, and consider the master equation

$$\frac{dP(N_j, t)}{dt} = \sum_k \sum_{l=k\pm 1} (E_k^+ E_l^- - 1) \mathcal{T}_{N_k-1, N_l+1 | N_k, N_l} P(N_j, t), \quad (8)$$

where $P(N_j, t)$ denotes the probability density function that there are N_j particles in compartment j at time t , for $j = 1, \dots, k_{max}$. Furthermore, $\mathcal{T}_{N_k-1, N_l+1 | N_k, N_l}$ denotes the transition rate from a state in which there are N_k, N_l particles in compartments k and l , respectively, to one in which there are $N_k - 1$ and $N_l + 1$ particles in those compartments, while particle numbers in all other compartments remain constant. Finally, the operators E_k^\pm are

defined by increasing (decreasing) N_k by one. A simple random walk model is defined by the transition rates

$$\mathcal{T}_{N_k-1, N_{k\pm 1}+1|N_k, N_{k\pm 1}} = \frac{D}{h^2} N_k. \quad (9)$$

Here, D is the diffusion coefficient that appears in equation (7). Henceforth, we refer to the model with transition rates (9) as the exact model of stochastic diffusion. As the transition rate is proportional to the number of particles in the outgoing box k , this indicates that the random walkers do not interact. The mean field equations associated with this stochastic model are given by

$$\frac{\partial \overline{N_k}}{\partial t} = \frac{D}{h^2} (\overline{N_{k+1}} - 2\overline{N_k} + \overline{N_{k-1}}). \quad (10)$$

Equation (10) can be viewed as an explicit finite-difference approximation of the diffusion equation (7) if we identify the densities with the particle numbers via $N_k(t) = hn(x, t)$, $x = kh$. We now propose the following, alternative transition rates for a model of stochastic diffusion:

$$\mathcal{T}_{N_k-1, N_{k\pm 1}+1|N_k, N_{k\pm 1}} = \frac{D}{h^2} \max(N_k - N_{k\pm 1}, 0). \quad (11)$$

Here, a random walker moves from compartment k to compartment $k \pm 1$ only if the number of particles in the outgoing box k exceeds the number in the incoming box $k \pm 1$. Transition rates (11) give rise to the same mean field equations (10) as transition rates (9) and, hence, also lead to the diffusion equation in the continuum limit. Indeed, the net average flux between neighboring compartments implied by transition rates (11) is the same as that implied by (9). However, we can expect that the noise associated with (9) will be larger, since the transition rates (11) are bounded above by the transition rates (9). In situations where noise changes the system dynamics, (11) might not be suitable. For this reason, we consider a combination of (9) and (11). We introduce a non-negative threshold Θ and the conditional difference model

$$\mathcal{T}_{N_k-1, N_{k\pm 1}+1|N_k, N_{k\pm 1}} = \begin{cases} \frac{D}{h^2} \max(N_k - N_{k\pm 1}, 0) & \text{if } N_k, N_{k\pm 1} \geq \Theta, \\ \frac{D}{h^2} N_k & \text{otherwise.} \end{cases} \quad (12)$$

We motivate (12) by noting that (9) represents the exact model of the random walkers, and should be used when particle numbers are low, since typically then relative noise levels are

significant. When particle numbers are high, i.e. $N_k, N_{k\pm 1} \geq \Theta$, we approximate (9) by (11), as for high particle numbers relative noise levels are typically low. We note that the mean of (11) and (9) are both given by (10), and Θ should be scaled with the lattice constant h . In particular, when the lattice constant is small, even if the total number of particles present in the system is large, the number of particles per compartment may not be large. In such cases the threshold would need to be reduced, but with caution since the application of (11) might lead to restrictive suppression of fluctuations.

We now aim to identify situations for which transition rates (9), (11) and (12) yield similar behaviour, and when they do not.

A. First-Passage Time problems

In many transport problems, it is important to know when a particle has first reached a certain site. We thus begin by investigating a first-passage time problem. We also refer the reader to ³³ for more studies of first-passage time problems. We suppose that initially N^0 particles are located at the left boundary compartment $k = 1$, and all other compartments are empty. We impose Dirichlet boundary conditions, so that $N_1 = N^0$ and $N_{k_{max}} = 0$ at all times, and the lattice constant is $h = 1$. We define the first-passage time at site k to be the time at which a particle first reaches compartment k . In Figure 2 we compare the average first-passage times obtained by averaging 1024 realisations of the exact random walk model (9), which coincides with the conditional difference model when $\Theta \rightarrow \infty$, and the random walk based on the difference model (11), which coincides with the conditional difference model when $\Theta = 0$. As expected, there is excellent agreement between the two models, because the first-passage times represent the time at which a single particle enters an empty compartment, and the transition rates for entry into an empty compartment are identical for (9) and (11).

When studying problems such as a reaction-diffusion system in which two, or more, particles of the same species need to be present at the same time and place, we should also establish whether changes in the transition rates of the random walk affect the time until several particles reach the same location. In such situations, we are concerned with the distribution of first times at which two particles are present in the same compartment. The first-passage times for two particles are shown in Figure 3. We notice that the difference




FIG. 2: First-passage times for $N^0 = 1000$ random walkers starting at compartment $k = 1$ at time 0, with $D = 10$. We present the average first-passage time for conventional random walkers based on the exact model defined by (9) (solid line) and for the random walkers based on the difference model (11) (dashed line). The results are averaged of 1024 simulations and are in excellent agreement.

model (11) predicts first passage times which are slightly longer than those predicted by the conditional difference model (12) with a threshold $\Theta = 5$, whose first-passage times are indistinguishable from those of the exact random walk model (9). This is because when one particle is present in a compartment, the rate at which a second particle enters that compartment is lower for the difference model, but identical for the conditional difference model when $\Theta \geq 2$ and the exact random walk model. Hence, when the system of interest involves reactions with two or more particles, the threshold needs to be set suitably high. In supplementary Figure S2, we present histograms of the number of particles in a given box at a given time, for the same scenario as above, comparing the distributions associated with different thresholds to those associated with the exact model $\Theta = \infty$. We find that, in general, lower thresholds lead to narrower distributions centered around the same mean, and that as the threshold levels are increases, the distribution approaches that obtained from the exact model. In supplementary Figure S3 we confirm that, as the threshold increases, the standard deviations of the exact model is recovered in a switch-like manner.

In Table I we show the average number of stochastic events needed to ensure that all




FIG. 3: The first time when a compartment k contains two particles is shown as a function of k for the exact random walk (9)(solid line, $\Theta = \infty$), the random walk based on the difference model (11)(dotted line, $\Theta = 0$) as well as the conditional difference model (12) (dashed line, $\Theta = 5$). Each plot is obtained by averaging 1024 simulations, with $D = 10$ and $N^0 = 1000$ random walkers at time 0 in compartment $k = 1$. The graphs for the exact random walk and the conditional difference model with $\Theta = 5$ coincide.

lattice sites have been visited at least once by at least one particle. We fix $D = 1$ and average over 1024 simulations. We see that the random walk difference model (11) needs only 4% of the number of events that the conventional random walk model (9) needs, with the conditional difference model (12) offering similar performance gains, although these improve as the threshold Θ is reduced. We remark also that the time to simulate a stochastic model with the conventional Gillespie algorithm is proportional to the total number of events that occur, since we repeat the same steps when simulating a single event. We conclude that the conditional difference model preserves the essential stochastic features of the exact random walk model while being considerably faster to simulate.

Number of Gillespie Events		
	$N^0 = 100$	$N^0 = 1000$
Exact random walk (9)	1.21×10^6	7.97×10^6
Difference model (11)	4.68×10^4	3.04×10^5
Conditional difference model (12), $\Theta = 5$	7.35×10^4	3.23×10^5
Conditional difference model(12), $\Theta = 10$	1.06×10^5	4.88×10^5

TABLE I: Comparison of the average number of Gillespie events needed to simulate the first-passage time problem, shown in Figure 2, for the exact random walk model (9), the difference model (11) and the conditional difference model (12) with two different thresholds $\Theta = 5, 10$. The simulations were stopped when every lattice site had been visited at least once by at least one particle.

IV. REACTION-DIFFUSION MODELS

A. Stochastic Fisher-Kolmogorov Equation

We now investigate a stochastic version of the Fisher-Kolmogorov equation. This equation can be viewed as a spatial-resolved birth-death process, in which identical particles move in a diffusive manner, divide and/or annihilate. The stochastic model is defined by the master equation

$$\begin{aligned} \frac{dP(N_j, t)}{dt} = & \sum_k \sum_{l=k\pm 1} (E_k^+ E_l^- - 1) \mathcal{T}_{N_k-1, N_l+1|N_k, N_l} P(N_j, t) \\ & + \sum_k ((E_k^- - 1) \mathcal{T}_{N_k+1|N_k} + (E_k^{+1} - 1) \mathcal{T}_{N_k-1|N_k}) P(N_j, t), \end{aligned} \quad (13)$$

where the birth and death transition rates are defined by

$$\begin{aligned} \mathcal{T}_{N_k+1|N_k} &= \lambda N_k, \\ \mathcal{T}_{N_k-1|N_k} &= \frac{\lambda}{\Omega h} N_k (N_k - 1), \end{aligned} \quad (14)$$

and the random walk transition rates are defined by either (9), (11) or (12). In (14), λ is a growth rate and $h\Omega$ is a measure of the carrying capacity of a single compartment, and, as before, h denotes the lattice constant. Consequently, if N_k is larger than $h\Omega$, then the annihilation process dominates the creation process. If we assume, as for the case of

the homogeneous birth-death process, that the particle number is sufficiently large that a van-Kampen limit applies, then this stochastic model results in mean field equations of the form

$$\frac{\partial \overline{N}_k}{\partial t} = \frac{D}{h^2} (\overline{N}_{k+1} - 2\overline{N}_k + \overline{N}_{k-1}) + \lambda \overline{N}_k \left(1 - \frac{\overline{N}_k}{\Omega h}\right). \quad (15)$$

Then, setting $n(x, t)h = N_k(t)$, $x = kh$, as before, we obtain the Fisher-Kolmogorov equation in the continuum limit:

$$\frac{\partial n}{\partial t} = D \frac{\partial^2 n}{\partial x^2} + \lambda n \left(1 - \frac{n}{\Omega}\right) \quad (16)$$

An interesting feature of equation (16) is that it generates stable travelling wave solutions which propagate with a constant speed of $2\sqrt{D\lambda}$ (see supplementary Figure S4 for an example of a travelling wave obtained from equation (16)). The wave speed of the stochastic Fisher-Kolmogorov equation differs from that of the corresponding PDE³⁴. Hence, the challenge for our algorithm is to reproduce this shift in wave speed, while being faster to simulate than the exact model. To this end, we now investigate travelling wave propagation in the stochastic Fisher-Kolmogorov equation, when the random walk part is modelled by the exact model (9) and by the difference model (11). For clarity, we do not apply our difference method to the reaction rates (14), since otherwise it will be less clear if the observed differences from using the difference model are due to the reaction of the diffusion parts of the model. Initial conditions are chosen so that the travelling wave front resembles that from the PDE. We impose Dirichlet boundary conditions so that $N_1 = \Omega$ and $N_{k_{max}} = 0$ for all times. We focus our analysis on the propagation of fully formed travelling wave solutions whose leading edge is at a distance from the boundary, so that the initial and boundary conditions do not influence the propagating wave. In Figure 4 we present snapshots of the average travelling wave profile from 256 different realisations of the stochastic Fisher-Kolmogorov equation (13). In Figure 4(a) the random walk is modelled with the exact model, (9), whereas in Figure 4(b) it is based on the difference model, (11). The results appear indistinguishable, indicating that (11) can reproduce the mean behaviour of (9). On closer inspection of the standard deviations, shown by the bars, we note that at the wave front (that is, in the regions where the solution involves small, but non-zero N_k), Figures 4(a) and 4(b) look similar, whereas away from the wave front (that is, in the regions where the solution fluctuates around $h\Omega = 100$), 4(b) has significantly less noise around the mean




FIG. 4: Instantaneous travelling wave profile in the stochastic Fisher-Kolmogorov equation (13) with parameters $D = 10, \lambda = 1, h = 1, k_{max} = 200$ and $\Omega = 100$. The random walk part of the model used for (a) is based on (9), whereas for (b) it is based on (11). The initial condition in each case was a traveling wave approximation and the plots shown were taken after a sufficiently large time to allow a stable travelling wave to form in each case.

Each plot shows the average of 256 different realisations, and the bars indicate the corresponding standard deviation.

than 4(a). The lower noise levels do not appear to have a significant effect on the shape or speed of the travelling wave. For a quantitative comparison of the two models, we plot the travelling wave speeds of the exact model (9) as well as the conditional difference model (12) in Figure 5. This is also a good testing ground for our algorithm, as the wave speed is known to be modified by stochastic effects³⁴. Here, we have obtained the wave speed in the same way as in¹⁶, averaging over 256 different simulations in each case, and measuring the speed at 15 different time intervals. We see that for low carrying capacities ($\Omega = 100$), the wave speed obtained by using the random walk difference model (11) ($\Theta = 0$) differs by approximately 1% from the travelling wave speed obtained by the conventional random walk (9), whereas the wave speed from the conditional difference model (12) differs by less than 1%. Furthermore, increasing the carrying capacity improves the agreement between the models. This is expected, as all three models converge to the mean field model given by equation (15) in the limit as $\Omega \rightarrow \infty$, and the mean field model has a wave speed close to the continuum model wave speed of $2\sqrt{D\lambda}$ ¹⁶.


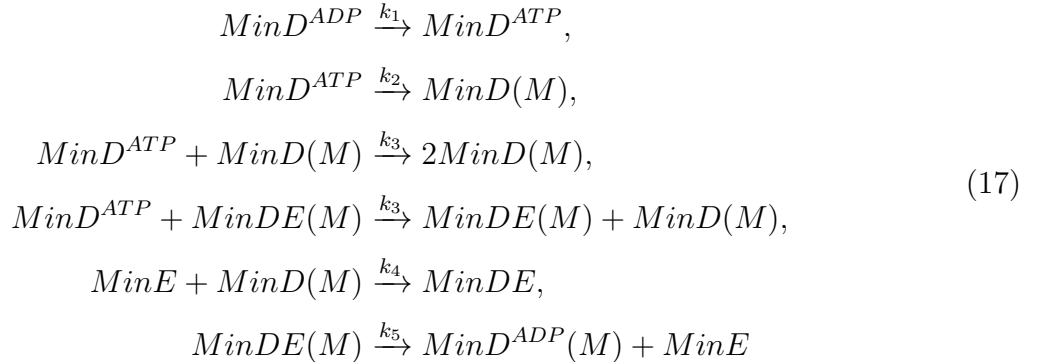


FIG. 5: Mean travelling wave speed for the stochastic Fisher-Kolmogorov equation (13) when $D = 1$, $\lambda = 1$, $k_{max} = 200$ and of the carrying capacity Ω . The solid line shows the wave speed for the conventional random walk model (9), the dotted line the difference model (11) ($\Theta = 0$) and the dashed line the conditional difference model (12) ($\Theta = 5$). Each plot was obtained by averaging over 256 realisations and measuring the speed at 5 time intervals for each simulation.

B. Min Oscillations in E.coli

We now investigate a model describing the interactions of Min proteins in E.coli. The deterministic PDE model was originally described in ³⁵, and stochastic versions were studied in ^{36,37}. We focus on a one-dimensional version of this model. The model describes reactions and diffusion of the molecules *MinD* and *MinE* which can exist in several states. $MinD^{ADP}$ and $MinD^{ATP}$ denote *MinD* sequestered in the cytosol and bound to ADP or ATP, respectively, whereas the sequestered form of *MinE* is simply denoted by *MinE*. The membrane-bound form of *MinD* is denoted by $MinD(M)$, and $MinDE(M)$ denotes the membrane-bound form of a complex of *MinD* and *MinE*. The reactions are defined by:



Furthermore, the cytosolic species $MinD^{ADP}$, $MinD^{ATP}$ and $MinE$ are able to diffuse through the cytosol with a macroscopic diffusion constant $D = 2.5\mu m^2/s$, which translates, as before, into rates on our lattice given by $\frac{D}{h^2}$. We assume diffusion on the membrane of $MinD(M)$ and $MinDE(M)$ is negligible.

The 1D model is obtained from the 3D model on a cylinder of length $L = 4\mu m$ and radius $R = 0.5\mu m$, by assuming diffusion occurs only along the main axis. We discretize along the length of the cylinder into cross-sectional disk-shaped compartments of length $h = 0.1\mu m$. Then, the reaction rates of the stochastic reaction, based on the PDE model³⁵, are given by $k_1 = \sigma_D^{ADP \rightarrow ATP} = 1s^{-1}$, $k_2 = \frac{2\sigma_D}{R} = 0.1s^{-1}$, $k_3 = \frac{\sigma_{dD}}{\pi R^2 h} = 0.019s^{-1}$, $k_4 = \frac{\sigma_E}{\pi R^2 h} = 1.18s^{-1}$, $k_5 = \sigma_{de} = 0.7s^{-1}$. Here, we use the original parameters σ from,³⁵ to estimate the reaction rates. We compare simulations of this system, when diffusion is described by the conditional difference model (12), with others where diffusion is modelled by the exact random walk model (9).

Supplementary Figure S5 shows the space-time evolution of the number of cytosolic MinE molecules with an initial total number of 1400 MinE and 6700 MinD proteins, for three different threshold values $\Theta = 0, 1, 10$, and the case of the exact random walk model (9), which corresponds to the conditional difference model with a large threshold. All four cases yield qualitatively similar results, with stable oscillations appearing shortly before $t = 100s$, with a period of approximately $70s$. However, noise levels are higher for the exact model, and when $\Theta = 10$, than when $\Theta = 0, 1$. We then quantified the distribution of molecular numbers. Figure 6 shows the distribution of four different molecular species, in the center of the simulation domain, corresponding to compartment $k = 20$, at time $t = 300$ and for 128 simulations. We see that $MinD^{ADP}$ and $MinD^{ATP}$ numbers (Figures 6(a), 6(b)) are unimodally distributed around a peak, and the conditional difference model with $\Theta = 1$ yields a narrower distribution than the exact model. By contrast, the distributions for $MinE$ and $MinD(M)$ (Figures 6(c), 6(d)) appear to decay monotonically with a peak at 0, and we observe no significant deviations between the exact model and the conditional difference model. To obtain more significant data, we average the standard deviations of those distributions over 200 time measurements (from $t = 300s$ to $t = 500s$, with measurements every $1s$ to avoid transient effects associated with the formation of the oscillations), in each compartment, and from each of the 128 different simulations. Figure 7 shows these standard deviations for five threshold values ($\Theta = 0, 1, 2, 10, 100000$), where the last case coincides with the exact




FIG. 6: Distribution of the number of Min molecules with an initial number of 1400 MinE and 6700 MinD proteins. We compare simulations of the stochastic model where diffusion is modeled by the conditional difference model, equation (12), with a threshold of $\Theta = 1$, and the case of the exact model, which is here obtained by choosing a threshold of $\Theta = 100000$. The histograms are obtained from the molecular numbers in a central compartment, $k = 20$, at time $t = 300s$, from 128 simulations of the model.

model. We see that for *MinE* and *MinD(M)*, there is excellent agreement between the different thresholds, whereas the standard deviations of the conditional difference models are slightly lower than those for the exact model for *MinD^{ADP}* and *MinD^{ATP}*. However, in those cases, relative noise is small as particle numbers are, on average, relatively large (see supplementary Figure S6, which also confirms that the means are almost identical for all chosen thresholds). Focusing on performance, we note that the computational cost decreases dramatically as Θ decreases. Figure 8 shows that absolute simulation time when $\Theta = 10$ decreased about 7 compared to the exact model, with a further 2 fold gain for thresholds




FIG. 7: Standard deviations from the distributions shown in 6. We have performed 128 simulations, and averaged the standard deviations over each space and time point, from $t = 300s$ until $t = 500s$.

of $\Theta = 0, 1, 2$. Based on these results, we conclude that choosing a threshold of $\Theta = 10$ offers excellent performance gains while maintaining the stochastic fluctuations associated with the exact model.

We stress that performance gains and quantitative similarity are parameter dependent. In supplementary Figure S7, we present results for the same scenario as before, but with only 25% of the particles present. We note that the standard deviations, similar to those shown in Figure 7, are similar for different thresholds. Corresponding performance gains are shown in Figure S8. The gains are small for $\Theta = 10$, but for $\Theta = 0, 1, 2$ the performance




FIG. 8: Simulation time on a 3 GHz Xeon machine of the simulations shown in Figures 6 and 7.

gains range from 13 to 25 fold. Hence, in this case, fixing $\Theta = 2$ may be preferred, as it offers great performance gains, but the behaviour of the species with the highest levels of relative fluctuation, $MinE$, behaves in a similar manner to the case of the exact random walk model.

V. CONCLUSIONS

In this paper, we have presented an algorithm which alters the transition rates that appear in reactions of the chemical master equation so that when two reactions have opposite effects on the state of the system, and when the particle number exceeds a threshold value, we replace the two reactions by a single reaction which represents the net effect of these two reactions. The mean behaviour of the stochastic model is preserved by this change of transition rates, and by setting the threshold sufficiently high it is possible to preserve essential stochastic features of the original model, such as variances, first-passage time in random walk models, or the stochastically modified wave speed in the stochastic Fisher-Kolmogorov equation.

We have applied the algorithm to diffusion and reaction-diffusion problems as diffusion often generates many practically unobservable events such as the swapping of positions of two identical particles, and such events are easily eliminated by our algorithm. Furthermore, as diffusion rates can often be considerably higher than reaction rates our algorithm can

help to significantly speed up simulations of the stochastic reaction-diffusion systems while preserving essential stochastic features. We have confirmed this for the problem of first-passage time, the modification of the wave speed of the Fisher-Kolmogorov system due to stochastic effects, and the distribution of particles in low copy numbers in the oscillatory system of Min proteins. Our current focus was on one-dimensional models. Future work will include a more detailed investigation of lattice dependencies, which can have an important effect on the behaviour of stochastic reaction-diffusion systems^{38–40}, or generalizations to more complex geometries as carried out in^{41,42}. There is also a need to test our method on simple lattices in higher dimensions. While the number of Gillespie events will be reduced in a similar manner as in one spatial dimension, events in the conditional difference model will lead to shifts in the total propensity, whereas movements in the bulk of the exact random walk model will only lead to local changes of propensities. This might make optimisations of the Gillespie algorithm in terms of efficient calculation of propensities potentially difficult.

A main advantage of our method for increasing the simulation speed of stochastic reaction or reaction-diffusion models is its simplicity of implementation. Unlike other algorithms designed to improve the simulation speed of such models, such as^{12–20}, ours can be implemented by simply redefining the transition rates in the Gillespie algorithm. Furthermore, it does not depend on the actual implementation used to simulate the stochastic process. Whilst we have used the Gillespie algorithm in this paper, it is straightforward to combine our algorithm with other solution methods of the master equation including the algorithms mentioned above.

ACKNOWLEDGMENTS

This publication was based on work supported in part by Award No KUK-C1-013-04, made by King Abdullah University of Science and Technology (KAUST). We are grateful to R. Erban, M. Flegg, A. McKane and M. Robinson for helpful discussions, and the anonymous referees for their helpful suggestions.

REFERENCES

- ¹N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry*, Vol. 1 (North holland, 1992).
- ²C. Gardiner, *Stochastic Methods* (Springer, 2010).
- ³D. T. Gillespie, A. Hellander, and L. R. Petzold, *The Journal of Chemical Physics* **138**, 170901 (2013).
- ⁴R. de la Cruz, P. Guerrero, F. Spill, and T. Alarcón, *The Journal of Chemical Physics* **143**, 074105 (2015).
- ⁵U. Dobramysl, S. Rüdiger, and R. Erban, arXiv preprint arXiv:1504.00146 (2015).
- ⁶A. J. Black and A. J. McKane, *Trends in Ecology & Evolution* **27**, 337 (2012).
- ⁷F. Spill, P. Guerrero, T. Alarcon, P. K. Maini, and H. M. Byrne, *Journal of Mathematical Biology* **70**, 485 (2014).
- ⁸M. Sturrock, A. Hellander, S. Aldakheel, L. Petzold, and M. A. Chaplain, *Bulletin of Mathematical Biology* **76**, 766 (2014).
- ⁹P. Guerrero, H. M. Byrne, P. K. Maini, and T. Alarcón, *Journal of Mathematical Biology*, 1 (2015).
- ¹⁰D. T. Gillespie, *Journal of Computational Physics* **22**, 403 (1976).
- ¹¹D. T. Gillespie, *The Journal of Physical Chemistry* **81**, 2340 (1977).
- ¹²M. A. Gibson and J. Bruck, *J. Phys. Chem. A* **104**, 1876 (2000).
- ¹³H.-P. Breuer, W. Huber, and F. Petruccione, *EPL (Europhysics Letters)* **30**, 69 (1995).
- ¹⁴A. Hellander and P. Lötstedt, *Journal of Computational Physics* **227**, 100 (2007).
- ¹⁵T. A. Henzinger, L. Mikeev, M. Mateescu, and V. Wolf, in *Proceedings of the 8th International Conference on Computational Methods in Systems Biology* (ACM, 2010) pp. 55–65.
- ¹⁶F. Spill, P. Guerrero, T. Alarcon, P. K. Maini, and H. Byrne, *Journal of Computational Physics* **299**, 429 (2015).
- ¹⁷D. T. Gillespie, *The Journal of Chemical Physics* **115**, 1716 (2001).
- ¹⁸D. T. Gillespie and L. R. Petzold, *The Journal of Chemical Physics* **119**, 8229 (2003).
- ¹⁹D. T. Gillespie, *Annu. Rev. Phys. Chem.* **58**, 35 (2007).
- ²⁰Y. Cao, D. T. Gillespie, and L. R. Petzold, *The Journal of Chemical Physics* **124**, 044109 (2006).

- ²¹J. Fu, S. Wu, and L. R. Petzold, *Journal of Computational Physics* **235**, 446 (2013).
- ²²S. Wu, J. Fu, and L. R. Petzold, *The Journal of Chemical Physics* **142**, 204108 (2015).
- ²³E. Flekkøy, J. Feder, and G. Wagner, *Physical Review E* **64**, 066302 (2001).
- ²⁴F. J. Alexander, A. L. Garcia, and D. M. Tartakovsky, *Journal of Computational Physics* **182**, 47 (2002).
- ²⁵F. J. Alexander, A. L. Garcia, and D. M. Tartakovsky, *Journal of Computational Physics* **207**, 769 (2005).
- ²⁶C. Li, U. Ebert, and W. Hundsdorfer, *Journal of Computational Physics* **231**, 1020 (2012).
- ²⁷M. B. Flegg, S. J. Chapman, and R. Erban, *Journal of The Royal Society Interface* (2011), 10.1098/rsif.2011.0574.
- ²⁸B. Franz, M. B. Flegg, S. J. Chapman, and R. Erban, *SIAM J. Appl. Math.* **73**, 1224 (2013).
- ²⁹M. B. Flegg, S. Hellander, and R. Erban, *Journal of Computational Physics* **289**, 1 (2015).
- ³⁰M. Robinson, M. Flegg, and R. Erban, *The Journal of Chemical Physics* **140**, 124109 (2014).
- ³¹K. R. Sanft and H. G. Othmer, *arXiv preprint arXiv:1503.05832* (2015).
- ³²C. A. Yates and M. B. Flegg, *Journal of The Royal Society Interface* **12** (2015), 10.1098/rsif.2015.0141.
- ³³A. Bezzola, B. B. Bales, R. C. Alkire, and L. R. Petzold, *Journal of Computational Physics* **256**, 183 (2014).
- ³⁴H.-P. Breuer, W. Huber, and F. Petruccione, *Physica D: Nonlinear Phenomena* **73**, 259 (1994).
- ³⁵K. C. Huang, Y. Meir, and N. S. Wingreen, *Proceedings of the National Academy of Sciences* **100**, 12724 (2003).
- ³⁶R. A. Kerr, H. Levine, T. J. Sejnowski, and W.-J. Rappel, *Proceedings of the National Academy of Sciences of the United States of America* **103**, 347 (2006).
- ³⁷D. Fange and J. Elf, *PLoS Comput Biol* **2**, e80 (2006).
- ³⁸S. A. Isaacson, *J. Chem. Phys.* **139**, 054101 (2013).
- ³⁹D. T. Gillespie, L. R. Petzold, and E. Seitaridou, *The Journal of Chemical Physics* **140**, 054111 (2014).
- ⁴⁰S. Hellander, A. Hellander, and L. Petzold, *Physical Review E* **91**, 023312 (2015).
- ⁴¹S. A. Isaacson and C. S. Peskin, *SIAM Journal on Scientific Computing* **28**, 47 (2006).

⁴²B. Drawert, S. Engblom, and A. Hellander, BMC Systems Biology **6**, 76 (2012).

VI. SUPPLEMENTARY INFORMATION

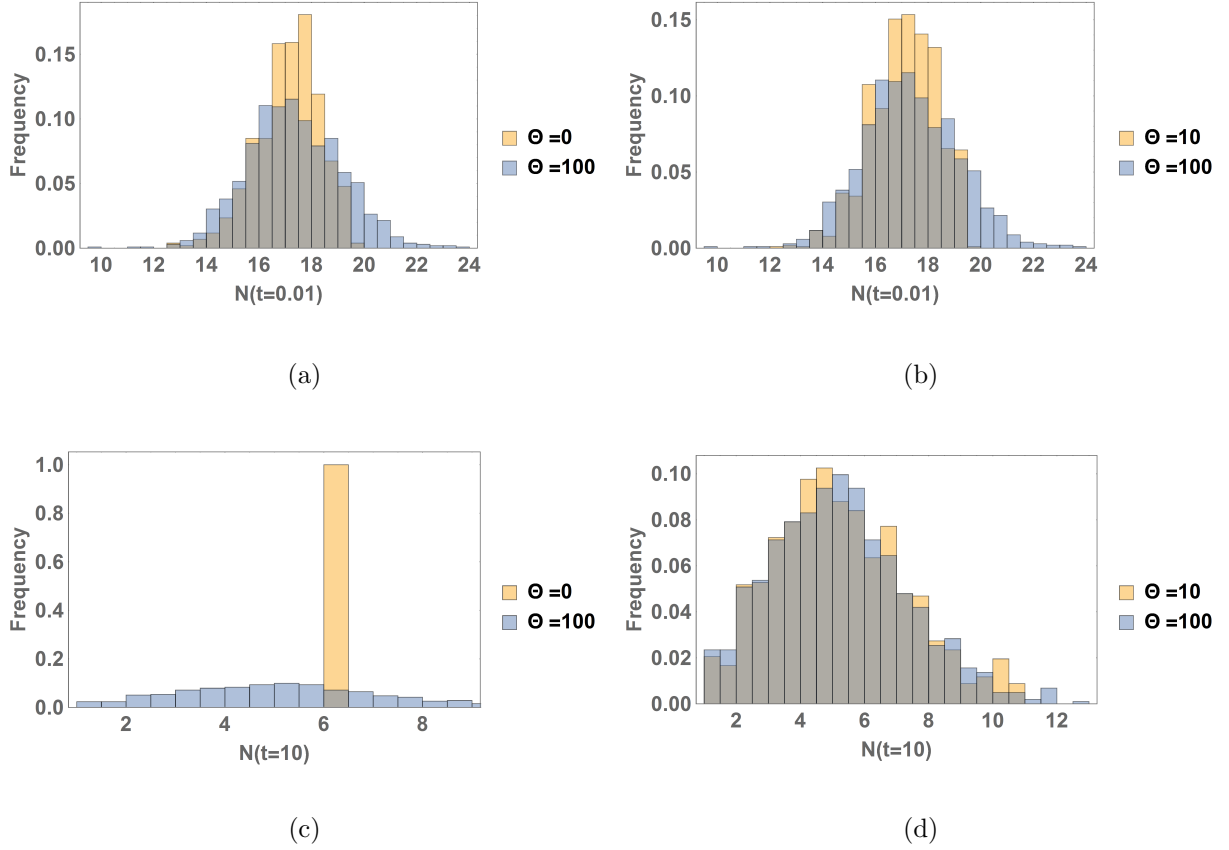





FIG. S1: Distribution of the birth-death process for the same parameters as in Figure 1, obtained from running 1024 different simulations, for two thresholds $\Theta = 0, 10$. In each case, we compare to the exact model, obtained by choosing a threshold $\Theta = 100$. The initial condition was $N(t = 0) = 20$, and the steady state is at $N = 6$. We notice that at $t = 0.01$, the distributions of all three cases, $\Theta = 0, 10, 100$ are similar, but, by construction, for $\Theta < 20$, no birth can occur so the distribution is bounded by $N(t = 0)$. At $t = 10$, the distributions for $\Theta = 10, 100$ look indistinguishable, whereas the distribution for $\Theta = 0$ goes to the delta distribution, and the solution is trapped at the absorbing state.




(a)



(b)



(c)



(d)

FIG. S2: Distribution of number of particles in compartment $k = 50$ for the diffusion process described in section III A, where initially 1000 particles are seeded in the compartment on the left, with all other compartments empty. After $t = 15$, we see that, while the conditional difference model with a threshold of $\Theta = 10$ closely matches the exact model ($\Theta = inf$), panel (b), the model with a threshold of $\Theta = 0$ shows a much narrower distribution than the exact model (panel (a)). At time $t = 20$, the case of $\Theta = 10$ now also deviates visibly from the exact case (panel (c)), whereas the case of $\Theta = 20$ offers close agreement (panel (d)).




FIG. S3: Standard Deviation of the diffusion process described in section III A, obtained from distributions similar to those shown in Figure S2. Here, measurements are taken at fixed time $t = 5$ in compartment $k = 30$, from 1024 simulations. We see that the standard deviations approach the ones of the exact model for thresholds around $\Theta = 10$.




FIG. S4: Travelling waves of the deterministic Fisher-Kolmogorov equation for parameters $D = 1$, $\lambda = 1$ and $\Omega = 100$. We see that a traveling wave is spreading with constant velocity $v = 2\sqrt{D\lambda} = 2$.







FIG. S5: Space-Time evolution of the number of MinE molecules with an initial number of 1400 MinE and 6700 MinD proteins. We compare simulations of the stochastic model where diffusion is modeled by the hybrid approach, equation (12), with different thresholds. Stable oscillations are forming after less than 100s. All four cases appear similar, but the cases with low thresholds, $\Theta = 0, 2$, appear less noisy than the case $\Theta = 10$ and the case of normal diffusion.




(a)



(b)




(c)




(d)

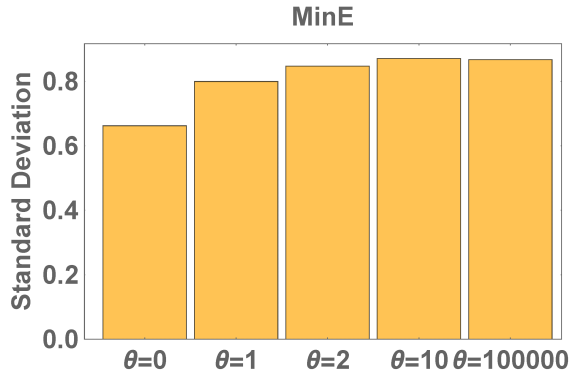
FIG. S6: Mean of the distributions shown in 6, corresponding to the standard deviations shown in Figure 7.




(a)



(b)



(c)



(d)

FIG. S7: Standard deviations from the distributions of particles, as in Figure 7, but for $\frac{1}{4}$ of the number of particles.




FIG. S8: Simulation performance as in Figure 8, but with only a quarter of the total number of particles.